

PIGGYBACKING VPN INFORMATION IN BGP FOR  
NETWORK BASED VPN ARCHITECTURES

This application claims the benefit of U.S.

5 Provisional Application No. 60/218,251, filed July 14, 2000.

FIELD OF THE INVENTION

This invention relates to communication systems  
implementing network based virtual private networks, and more  
particularly to distribution within such systems of information  
10 relating to the virtual private networks.

BACKGROUND OF THE INVENTION

In communication systems consisting of multiple  
Autonomous Systems (ASs), a routing protocol is required  
whereby the ASs can share information about sites and hosts  
15 that can be reached within each ASs. A common routing  
protocol, and the routing protocol used in the Internet for  
Internet Protocol version 4 (IPv4) traffic, is the Border  
Gateway Protocol version 4 (BGP-4). BGP-4 is defined in  
Rekhter, Y., and Li, T., "A Border Gateway Protocol 4 (BGP-4)",  
20 IETF RFC1771, T. J. Watson Research Center, cisco, March 1995,  
incorporated by reference herein. BGP-4 allows Border Gateway  
Protocol Speakers (BGP Speakers) to automatically and  
periodically exchange Network Layer Reachability Information  
(NLRI). NLRI is sent as part of an Update message between  
25 border gateways. Each Update message may advertise one  
feasible route. A feasible route includes a NLRI field,  
composed of a list of Internet Protocol (IP) addresses which  
are destinations of the route, and zero or more Path  
Attributes, which each include an Attribute Flag, an Attribute  
30 Type Code and an Attribute Value. BGP-4 defines several  
Attribute Type Codes, but allows new Attribute Type Codes to be

defined. Using NLRI collected over time from Update messages, each BGP Speaker builds routing databases whereby a route to a destination within the communication system can be determined.

Due to the definition of some Attribute Value fields within BGP-4, BGP-4 is capable of carrying routing information only for IPv4 traffic. However some communication systems implement different networking systems. For example, different Open System Interconnect (OSI) layer-2 or layer-3 protocols may be used within the communication system. One solution is

proposed in Bates, T., Chandra, R., Katz, D., and Rekhter, Y., "Multiprotocol Extensions for BGP-4", IETF RFC2283, Cisco Systems, Juniper Networks, February 1998, incorporated by reference herein, in which extensions to BGP-4 are defined.

The resulting BGP Multiprotocol Extensions (BGP-MP) enable BGP-4 to distribute routing information for multiple network layer protocols. BGP-MP introduces two new Path Attributes. The first new Path Attribute, Multiprotocol Reachable NLRI (MP\_REACH\_NLRI), carries a set of destinations reachable through the BGP Speaker which sent the Update message, together with next hop information. The new Attribute Type Code of MP\_REACH\_NLRI is defined to have a value of "14".

MP\_REACH\_NLRI is defined as an optional Path Attribute using the Attribute Flag (to allow for backward compatibility). The Attribute Value of MP\_REACH\_NLRI consists of several fields.

The first field in the Attribute Value is an Address Family Identifier field, which identifies the network layer protocol associated with the destinations identified in an NLRI field. The second field in the Attribute Value is a Subsequent Address Family Identifier (SAFI) field, which provides additional information about the type of NLRI carried in the Path Attribute. The second new Path Attribute, Multiprotocol Unreachable NLRI (MP\_UNREACH\_NLRI) carries a set of routes which are to be withdrawn from the routing database of the BGP

Speaker which receives the Update message. The new Attribute Type Code of MP\_UNREACH\_NLRI is defined to have a value of "15". Like MP\_REACH\_NLRI, MP\_UNREACH\_NLRI is defined as an optional Path Attribute, and its Attribute Value includes an 5 Address Family Identifier field and a SAFI field. In summary, when a BGP Speaker receives an Update message having an Attribute Type Code of "14" or "15", the BGP Speaker realizes that the destinations listed in the NLRI field are associated with the network layer protocol identified by the Address 10 Family Identifier field.

Network based Virtual Private Networks (VPNs) are emulations of private wide area networks using public or other party backbone facilities. Network based VPNs allow a business to run a wide area network over a large geographic area without 15 having to purchase and manage its own backbone, thereby saving costs. For example, a business may own computers at two geographically separate sites, and wish to connect them as a network. Network based VPNs provide a means for the business to link the two sites using the backbone facilities of someone 20 else (referred to as a provider). Network based VPNs can be implemented using either a piggybacking model (Rosen, E., et al., "BGP/MPLS VPNs", IETF Internet-Draft update of RFC2547, May 2000, incorporated by reference herein, provides an example of a piggybacking model) or a Virtual Router model (see, for 25 example, Ould-Brahim, H., et al., "Network based IP VPNs using Virtual Routers", IETF Internet-Draft, July 2000, incorporated by reference herein). Current methods for implementing network based VPNs in either model require the same networking system to be used in each VPN as is used in the backbone. However, 30 when purchasing VPN services from a provider a business with an existing private network may not wish to alter its networking system in order to comply with the networking system used by the backbone and the other VPNS. There is a need for a VPN

information distribution scheme which allows different networking systems to be used by each VPN and by the backbone.

#### SUMMARY OF THE INVENTION

The present invention provides a Border Gateway  
5 Protocol Speaker (BGP Speaker) in a communication system which  
implements at least one network based Virtual Private Network  
(NB-VPN) across a backbone. Each NB-VPN uses an Open System  
Interconnect (OSI) layer-2 protocol and an OSI layer-3  
protocol, and at least one NB-VPN uses an OSI layer-2 protocol  
10 different from an OSI layer-2 protocol used by the backbone or  
uses an OSI layer-3 protocol different from an OSI layer-3  
protocol used by the backbone. The BGP Speaker transmits an  
Update message in conformance with a Border Gateway Protocol  
(BGP), and the Update message further includes, Virtual Private  
15 Network (VPN) Membership information, VPN Reachability  
information, and Tunnel Mechanism information. The Update  
message further includes a field indicating a topology of the  
NB-VPN and a field indicating a VPN Reachability Mode. In  
addition to the BGP Speaker, the invention provides the data  
20 format itself.

The VPN Membership information includes at least one  
VPN Identification (VPN-ID) field and a Number of VPN-IDs  
field. The VPN Reachability information includes zero or more  
VPN Reachability Entries. Each VPN Reachability Entry includes  
25 a VPN Reachability Type field, a Length field, and a VPN  
Reachability Value field. The Tunnel Mechanism information  
includes zero or more VPN Tunnel Entries, each VPN Tunnel Entry  
including a Tunnel Type field, a Length field and a Tunnel  
Value field.

30 The Update message includes a unique Subsequent  
Address Family Identifier (SAFI) value indicating that the

Update message includes the information relating to the NB-VPNs and the information relating to networking systems used by the NB-VPNs. A SAFI value of "129" is proposed.

The invention also provides a Virtual Router  
5 receiving an Update message being in conformance with a Border  
Gateway Protocol (BGP) and further including information  
relating to a NB-VPN to which the VR belongs and information  
relating to networking systems used by the NB-VPN to which the  
VR belongs, and the VR including instructions for establishing  
10 an OSI layer-2 connection to at least one other VR in the  
NB-VPN.

The method allows VPN routing information to be distributed within a communication system employing multiple OSI layer-2 or layer-3 protocols. In addition, a VR can discover OSI layer-2 information and establish a layer-2 link with another VR, thereby avoiding the need to establish a layer-3 link through the backbone.

Other aspects and features of the present invention will become apparent to those ordinarily skilled in the art upon review of the following description of specific embodiments of the invention in conjunction with the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention will now be described in greater detail  
25 with reference to the accompanying diagrams, in which:

FIG. 1 is a block diagram of a network based Virtual Private Network (VPN) using a virtual router model;

FIG. 2 is a block diagram of a network based VPN using a piggybacking model;

FIG. 3a is a chart of a format of Network Layer Reachability Information (NLRI);

FIG. 3b is a chart of a format of a VPN Reachability Entry within the NLRI; and

5 FIG. 3c is a chart of a format of a VPN Tunnel Entry within the NLRI.

#### DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Referring to FIG. 1, a network based Virtual Private Network (VPN) implemented using a Virtual Router (VR) model within a communication system is shown. Two private networks, a first private network being located at a first site 10 and a second private network being located at a second site 12, are operated by a customer who wishes to establish a VPN to connect both private networks. Hosts at the first site 10 are connected, either directly or indirectly via other hosts within the first site, to a first Customer Edge (CE) device 24. A CE device may be a router, a switch, or, if the private network to which the CE device belongs is a single host, a host. Hosts at the second site 12 are connected, either directly or 10 indirectly, to a second CE device 26. The first CE device 24 is connected to a first VR 14 over an access link, and the second CE device 26 is connected to a second VR 18 over another access link. The first VR 14 is located within a first Provider Edge (PE) device 16 and is connected to a first 15 Backbone Virtual Router (BVR) 15, also within the first PE device 16. A PE device may be a router or a switch. The second VR 18 is located within a second PE device 20 and is connected to a second BVR 19, also within the second PE device 20. Both VRs maintain routing tables which define a 20 reachability for the VPN. The BVRs 15 and 19 are connected to a backbone 22, which is not aware of the VPN. The backbone 22 25

is generally a shared network infrastructure. Neither the backbone 22 nor the PE devices 16 and 20 are operated by the customer. If a Border Gateway Protocol (BGP) is used by the BVRs to advertise routes and to allow the routing tables of the 5 VRs to be populated, then each BVR is a BGP Speaker, in that each BVR transmits Update messages in accordance with BGP.

More than one VPN may be present within the communication system, in which case any PE device may contain more than one VR, each VR maintaining routing tables for the 10 VPN to which it is connected, and each VR within a PE device being connected to the BVR in the PE device. Two VPNs (VPN1 and VPN2) are shown in FIG. 1. If there is more than one VPN, then there may be more than one backbone. All VRs belonging to the same VPN must have an identical VPN Identification (VPN-ID) 15 unique to the VPN, an example format of which is defined in Fox, B. et al, "Virtual Private Networks Identifier", IETF RFC2685, September 1999 (RFC2685), incorporated by reference herein.

A network based Virtual Private Network (VPN) can 20 also be implemented using a piggybacking model within a communication system, as shown in FIG. 2. Unlike the VR model, VRs are not used. Rather, the first CE device 24 is connected to the first PE device 16 over an access link, and the second CE device 26 is connected to the second PE device 20 over 25 another access link. Each PE device is connected to the backbone 22. Each PE device maintains routing tables which define a reachability for each VPN in each site to which the PE device is connected. Each PE device is a BGP Speaker, in that each PE device transmits Update messages in accordance with 30 BGP.

The backbone and each VPN employ Open System Interconnect (OSI) layer-2 protocols and OSI layer-3 protocols. The communication system is a multiservice system in that the VPNs need not employ the same OSI layer-2 protocol as that used by the backbone, and need not employ the same OSI layer-3 protocol as that used by the backbone. Furthermore, a VPN need not use the same OSI layer-2 or OSI layer-3 protocol as is used by other VPNs in the communication system.

BGP Speakers periodically transmit Update messages to each other using a Multiprotocol Extension to the Border Gateway Protocol (BGP-MP). An Update message may include a Multiprotocol Reachability NLRI (MP\_REACH\_NLRI) Path Attribute having an Attribute Type Code of "14". A MP\_REACH\_NLRI is used by the BGP Speaker that transmits the Update message to advertise a feasible route to destinations listed in the NLRI. The routing tables maintained by the BGP Speaker which receives the Update message may be updated to include the feasible route. An Update message may also include a Multiprotocol Unreachability NLRI (MP\_UNREACH\_NLRI) Path Attribute having an Attribute Type Code of "15". A MP\_UNREACH\_NLRI is used by the BGP Speaker that transmits the Update message to advertise routes which are no longer feasible, and which should be removed from the routing tables maintained by the BGP Speaker which receives the Update message. Both of the MP\_REACH\_NLRI and MP\_UNREACH\_NLRI Path Attributes may include a set of Network Layer Reachability Information (NLRI) fields, each NLRI field listing one or more destination addresses of a route. Both of these Path Attributes include an Address Family Identifier field, which identifies a network layer protocol associated with the destination addresses listed in the NLRI field. Both of these Path Attributes also include a Subsequent Address Family Identifier (SAFI) field, which provides additional information about the destination addresses listed

in the NLRI field. A unique value for the SAFI field of both MP\_REACH\_NLRI and MP\_UNREACH\_NLRI is used to indicate that the NLRI fields contain VPN information. A value of "129" may be used as the unique value for the SAFI field, as the value of 5 "129" is not currently in use and is not a reserved value. The VPN information contained in each NLRI field produced by a BGP Speaker is shown in FIG. 3a to 3c.

Referring first to FIG. 3a, a NLRI field includes a Length of NLRI field 100, which indicates a length of the NLRI 10 field in bits and is two octets in length. A NLRI field includes VPN Membership Information. VPN Membership

Information includes a Number of VPN-Identifications (VPN-IDs) 15 field 105 and at least one VPN-ID field 110. The Number of VPN-IDs field 105 indicates the number of VPN-ID fields in the NLRI, and is two octets in length. Each VPN-ID field 20 identifies a VPN to which the NLRI relates, and is preferably in conformance with the format specified by RFC2685. A NLRI field also includes a VPN Topology Attribute field 115, which is used to indicate a VPN topology and is two octets in length.

A NLRI field also includes VPN Reachability 25 Information which defines routes by which the VPNs to which the NLRI relates can be reached. The VPN Reachability Information includes a Length of VPN Reachability Entries field 125, and zero or more VPN Reachability Entries 30 130. The Length of VPN Reachability Entries field 125 indicates a total length of the VPN Reachability Entry fields 130 in bits, and is two octets in length. Each VPN Reachability Entry field 130 has a format as shown in FIG. 3b, and includes a VPN Reachability Type field 150, a Length field 155, and a VPN Reachability Value field 160. The VPN Reachability Type field 150 indicates a type of 35 VPN route, and is two octets in length. The Length field 155 indicates a length of the VPN Reachability Value field 160 in

bits, and is one octet in length. The VPN Reachability Value field 160 contains a VPN route, along with any information relevant to the VPN route, and has a variable length.

A NLRI field also includes Tunnel Mechanism

5 information. A tunnel allows opaque transport of VPN packets across the backbone, such that packet forwarding within the backbone is independent of VPN address fields within the packet. The Tunnel Mechanism information includes a Length of VPN Tunnel Entries field 135 and zero or more VPN Tunnel Entry  
10 fields 140. The Length of VPN Tunnel Entries field 135 indicates a total length of the VPN Tunnel Entry fields 140 in  
0 bits, and is two octets in length. Each VPN Tunnel Entry field  
0 has a format as shown in FIG. 3c, and includes a Tunnel Type  
0 field 165, a Length field 170, and a Tunnel Value field 175.  
15 The Tunnel Type field 165 indicates a type of tunnelling  
0 mechanism, and is two octets in length. The Length field 170  
0 indicates a length of the Tunnel Value field in bits, and is  
0 one octet in length. The Tunnel Value field 175 carries  
0 information related to an endpoint of the tunnel, and has a  
0 variable length. The Tunnel Value field 175 can carry, for  
20 example, address information, Quality of Service information,  
0 and tunnel mechanism parameters.

A NLRI field also includes a VPN Reachability Mode field 120. The VPN Reachability Mode field 120 indicates  
25 whether a piggybacking model or a VR model is being used by the  
VPNs to which the NLRI relates, and is one octet in length. A  
NLRI field may also include an Other VPN Information field 145  
for passing various VPN information. The Other VPN Information  
field 145 may carry vendor specific information, for example  
30 Quality of Service parameters or passwords. The Other VPN  
Information field 145 is optional in a NLRI field.

A network based VPN implemented using a VR architecture may also use the VPN information depicted in FIG. 3a to 3c to establish Open System Interconnect (OSI) layer-2 links between VRs. This removes the necessity of a layer-3 link within the backbone, providing additional security for traffic flowing between the VRs.

What has been described is merely illustrative of the application of the principles of the invention. Other arrangements and methods can be implemented by those skilled in the art without departing from the spirit and scope of the present invention. For example, the various field lengths and field values that have been described are proposed for use as a standard protocol. Different field lengths and different field values may be varied without departing from the scope of the invention. Formats other than that specified in RFC2685 can be used to describe the VPN-ID fields 110.

5  
10  
15